The Frobenius anatomy of word meanings I: subject and object relative pronouns

Mehrnoosh Sadrzadeh
Queen Mary University of London
School of Electronic Eng. and Computer Science
mehrs@eecs.qmul.ac.uk

Stephen Clark
University of Cambridge
Computer Laboratory
sc609@cam.ac.uk

Bob Coecke
University of Oxford
Dept. of Computer Science
coecke@cs.ox.ac.uk

arXiv:1404.5278v1 [cs.CL] 21 Apr 2014

Abstract. This paper develops a compositional vector-based semantics of subject and object relative pronouns within a categorical framework. Frobenius algebras are used to formalise the operations required to model the semantics of relative pronouns, including passing information between the relative clause and the modified noun phrase, as well as copying, combining, and discarding parts of the relative clause. We develop two instantiations of the abstract semantics, one based on a truth-theoretic approach and one based on corpus statistics.

1 Introduction

Ordered algebraic structures and sequent calculi have been used extensively in Computer Science and Mathematical Logic. They have also been used to formalise and reason about natural language. Lambek (1958) [17] used the ordered algebra of residuated monoids to model grammatical types, their juxtapositions and reductions. Relational words such as verbs have implicative types and are modelled using the residuals of the monoid multiplication [23,24]. Later, Lambek (1999) [18] simplified these algebraic calculi in favour of *pregroups*. Here, there are no binary residual operations, but each element of the algebra has a left and a right residual or adjoint.

In terms of semantics, pregroups do not naturally lend themselves to a model-theoretic treatment [22]. However, pregroups are suited to a radically different treatment of semantics, namely distributional semantics [27]. Distributional semantics uses vector spaces based on contextual co-occurrences to model the meanings of words. Coecke et al. (2010) [8] show how a *compositional* semantics can be developed within a vector-based framework, by exploiting the fact that vector spaces with linear maps and pregroups both have a compact closed categorical structure [15,25]. Some initial attempts at implementation include Grefenstette and Sadrzadeh (2011a) [10] and Grefenstette and Sadrzadeh (2011b) [11].

One problem with the distributional approach is that it is difficult to see how the meanings of some words — e.g. logical words such as *and*, *or*, and relative

pronouns such as *who, which, that, whose* — can be modelled contextually. Our focus in this paper is on relative pronouns in the distributional compositional setting.

The difficulty with pronouns is that the contexts in which they occur do not seem to provide a suitable representation of their meanings: pronouns tend to occur with a great many nouns and verbs. Hence, if one applies the contextual co-occurrence methods of distributional semantics to them, the result will be a set of dense vectors which do not discriminate between different meanings. The current state-of-the-art in compositional distributional semantics either adopts a simple method to obtain a vector for a sequence of words, such as adding or mutliplying the contextual vectors of the words [20], or, based on the grammatical structure, builds linear maps for some words and applies these to the vector representations of the other words in the string [2,10]. Neither of these approaches produce vectors which provide a good representation for the meanings of relative clauses.

In the grammar-based approach, one has to assign a linear map to the relative pronoun, for instance a map f as follows:

$$\overrightarrow{\text{men who like Mary}} = f(\overrightarrow{\text{men}}, \overrightarrow{\text{like Mary}})$$

However, it is not clear what this map should be. Ideally, we do not want it to depend on the frequency of the co-occurrence of the relative pronoun with the relevant basis vectors. But both of the above mentioned approaches rely heavily on the information provided by a corpus to build their linear maps. The work of Baroni and Zamparelli (2010) [2] uses linear regression and approximates the context vectors of phrases in which the target word has occurred, and the work of Grefenstette and Sadrzadeh (2011a) [10] uses the sum of Kronecker products of the arguments of the target word across the corpus.

The semantics we develop for relative pronouns and clauses uses the general operations of a Frobenius algebra over vector spaces [7] and the structural categorical morphisms of vector spaces. We do not rely on the co-occurrence frequencies of the pronouns in a corpus and only take into account the structural roles of the pronouns in the meaning of the clauses. The computations of the algebra and vector spaces are depicted using string diagrams [12], which depict the interactions that occur among the words of a sentence. In the particular case of relative clauses, they visualise the role played by the relative pronoun in passing information between the clause and the modified noun phrase, as well as copying, combining, and even discarding parts of the relative clause.

We develop two instantiations of the abstract semantics, one based on a truth-theoretic approach, and one based on corpus statistics, where for the latter the categorical operations are instantiated as matrix multiplication and vector

component-wise multiplication. As a result, we will obtain the following for the meaning of a subjective relative clause:

$$\overrightarrow{\text{men who like Mary}} = \overrightarrow{\text{men}} \odot (\overline{\text{love}} \times \overrightarrow{\text{Mary}})$$

The rest of the paper introduces the categorical framework, including the formal definitions relevant to the use of Frobenius algebras, and then shows how these structures can be used to model relative pronouns within the compositional vector-based setting.

2 Compact Closed Categories and Frobenius Algebras

We briefly review the main concepts of compact closed categories and Frobenius Algebras, focusing on the concepts and constructions that we are going to use in our models. For a formal presentation, see [15,16,1]; for an informal introduction, see [6].

A compact closed category has objects $A, B, C, ...$, morphisms $f\colon A \to B$, $g\colon B \to C$, composition $g \circ f$, a monoidal tensor $A \otimes B$ that has a unit I, and for each object A two objects A^r and A^l together with the following morphisms:

$$A \otimes A^r \xrightarrow{\epsilon^r_A} I \xrightarrow{\eta^r_A} A^r \otimes A$$
$$A^l \otimes A \xrightarrow{\epsilon^l_A} I \xrightarrow{\eta^l_A} A \otimes A^l$$

These structural morphisms of the compact closed structure satisfy the following equalities, some times referred to as *yanking*, where 1_A is the identity morphism on the object A:

$$(1_A \otimes \epsilon^l_A) \circ (\eta^l_A \otimes 1_A) = 1_A \qquad (\epsilon^r_A \otimes 1_A) \circ (1_A \otimes \eta^r_A) = 1_A$$
$$(\epsilon^l_A \otimes 1_A) \circ (1_{A^l} \otimes \eta^l_A) = 1_{A^l} \qquad (1_{A^r} \otimes \epsilon^r_A) \circ (\eta^r_A \otimes 1_{A^r}) = 1_{A^r}$$

A pregroup is a partially ordered compact closed category. This means that the objects are elements of a partially ordered monoid and between any two objects p, q there exists a morphism of type $p \to q$ iff $p \leq q$. Compositions of morphisms are obtained by transitivity and the identities by reflexivity of the partial order. The tensor of the category is the monoid multiplication, and the four yanking inequalities become as follows:

$$\epsilon^r_p = p \cdot p^r \leq 1 \qquad \epsilon^l_p = p^l \cdot p \leq 1$$
$$\eta^r_p = 1 \leq p^r \cdot p \qquad \eta^l_p = 1 \leq p \cdot p^l$$

In the rest of the paper, by $Preg$ we either mean any or a particular pregroup, as determined by the context.

Finite dimensional vector spaces and linear maps also constitute a compact closed category, to which we refer as $FVect$. Finite dimensional vector spaces $V, W, ...$ are objects of this category; linear maps $f\colon V \to W$ are its morphisms with composition being the composition of linear maps. The tensor product of the spaces $V \otimes W$ is the linear algebraic tensor product, whose unit is the scalar field of the vector spaces; in our case this is the field of reals $\mathbb{R}$. As opposed to the tensor product in $Preg$, the tensor between vector spaces is symmetric, hence we have a natural isomorphism $V \otimes W \cong W \otimes V$. As a result of the symmetry of the tensor, we have $V^l \cong V^r \cong V^*$, where V^* is the dual space of V. When the basis vectors of vector spaces are fixed, we furthermore obtain that $V^* \cong V$. Elements of a vector space V, that is vectors $\overrightarrow{v} \in V$, are presented by linear maps of type $\mathbb{R} \to V$, which can be defined by setting $1 \mapsto \overrightarrow{v}$ when relying on linearity. We still denote these maps by $\overrightarrow{v}$.

Given a basis $\{r_i\}_i$ for a vector spaces V, the epsilon maps are given by the inner product extended by linearity, that is we have:

$$\epsilon^l = \epsilon^r\colon V^* \otimes V \to \mathbb{R} \quad \text{given by} \quad \sum_{ij} c_{ij}\ \psi_i \otimes \phi_j \quad \mapsto \quad \sum_{ij} c_{ij} \langle \psi_i \mid \phi_j \rangle$$

Eta maps are defined as the linear map representing a vector in $V \otimes V^*$:

$$\eta^l = \eta^r\colon \mathbb{R} \to V \otimes V^* \quad \text{given by} \quad 1 \quad \mapsto \quad \sum_i r_i \otimes r_i$$

As mentioned above, the assignment of an image to $1 \in \mathbb{R}$ extends to all numbers by linearity.

A Frobenius algebra over a symmetric monoidal category $(\mathcal{C}, \otimes, I)$ is a tuple $(X, \Delta, \iota, \mu, \zeta)$ where for X an object of $\mathcal{C}$, the triple

$$(X, \Delta\colon X \to X \otimes X, \iota\colon X \to I)$$

is an internal comonoid, the triple

$$(X, \mu\colon X \otimes X \to X, \zeta\colon I \to X)$$

is an internal monoid, and which together moreover satisfy the following *Frobenius* condition:

$$(\mu \otimes 1_X) \circ (1_X \otimes \Delta) \ = \ \Delta \circ \mu \ = \ (1_X \otimes \mu) \circ (\Delta \otimes 1_X)$$

Frobenius algebras were originally introduced in 1903 by F. G. Frobenius in the context of proving representation theorems for group theory [9]. Since then,

they have found applications in other fields of mathematics and physics, e.g. in topological quantum field theory [16]. The above general categorical definition is due to Carboni and Walters [3].

In the category of finite dimensional vector spaces and linear maps **FVect**, any vector space V with a fixed basis $\{\overrightarrow{v_i}\}_i$ has a Frobenius algebra over it, explicitly given by:

$$\Delta :: \overrightarrow{v_i} \mapsto \overrightarrow{v_i} \otimes \overrightarrow{v_i} \qquad\qquad \iota :: \overrightarrow{v_i} \mapsto 1$$

$$\mu :: \overrightarrow{v_i} \otimes \overrightarrow{v_j} \mapsto \delta_{ij}\, \overrightarrow{v_i} := \begin{cases} \overrightarrow{v_i} & i = j \\ \overrightarrow{0} & i \neq j \end{cases} \qquad\qquad \zeta :: 1 \mapsto \textstyle\sum_i \overrightarrow{v_i}$$

Frobenius algebras over vector spaces with orthonormal bases are moreover *special* and *commutative*. A commutative Frobenius algebra satisfies the following two conditions for $\sigma : X \otimes Y \to Y \otimes X$ the symmetry morphism of $(\mathcal{C}, \otimes, I)$:

$$\sigma \circ \Delta = \Delta \qquad\qquad \mu \circ \sigma = \mu$$

A special Frobenius algebra is one that satisfies the following axiom:

$$\mu \circ \Delta = 1$$

The vector spaces of distributional models that we work with have fixed orthonormal bases, hence they have special commutative Frobenius algebras over them.

The comonoid's comultiplication of a special commutative Frobenius algebra over a vector space encodes vectors of lower dimensions into vectors of higher dimensional tensor spaces; this operation is usually referred to by *copying*. In linear algebraic terms, $\Delta(\overrightarrow{v}) \in V \otimes V$ is a diagonal matrix whose diagonal elements are the coefficients (or weights) of $\overrightarrow{v} \in V$. The corresponding algebraic operations encode elements of higher dimensional tensor spaces into lower dimensional spaces; this operation is usually referred to by *uncopying*. For $\overrightarrow{w} \in V \otimes V$, when represented as a matrix, we have that $\mu(\overrightarrow{w}) \in V$ is a vector consisting only of the diagonal elements of $\overrightarrow{w}$.

As a concrete example, take V to be a two dimensional space with basis $\{\overrightarrow{v_1}, \overrightarrow{v_2}\}$, then the basis of $V \otimes V$ is $\{\overrightarrow{v_1} \otimes \overrightarrow{v_1}, \overrightarrow{v_1} \otimes \overrightarrow{v_2}, \overrightarrow{v_2} \otimes \overrightarrow{v_1}, \overrightarrow{v_2} \otimes \overrightarrow{v_2}\}$. For a vector $v = a\overrightarrow{v_1} + b\overrightarrow{v_2}$ in V we have:

$$\Delta(v) = \Delta \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} = a\, \overrightarrow{v_1} \otimes \overrightarrow{v_1} + b\, \overrightarrow{v_2} \otimes \overrightarrow{v_2}$$

And for $\overrightarrow{w} = a\, \overrightarrow{v_1} \otimes \overrightarrow{v_1} + b\, \overrightarrow{v_1} \otimes \overrightarrow{v_2} + c\, \overrightarrow{v_2} \otimes \overrightarrow{v_1} + d\, \overrightarrow{v_2} \otimes \overrightarrow{v_2}$ in $V \otimes V$, we have:

$$\mu(w) = \mu \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} a \\ d \end{pmatrix} = a\, \overrightarrow{v_1} + d\, \overrightarrow{v_2}$$

3 String Diagrams

The framework of compact closed categories and Frobenius algebras comes with a diagrammatic calculus that visualises the derivations thereof. These diagrams also simplify the categorical and vector space computations of our setting to a great extent.

We briefly introduce the fragment of this calculus that we will use. Morphisms are depicted by boxes and objects by lines, representing their identity morphisms. For instance a morphism $f\colon A \to B$, and an object A with the identity arrow $1_A\colon A \to A$ are depicted as follows:

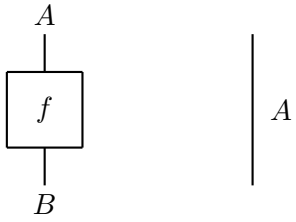

The tensor products of the objects and morphisms are depicted by juxtaposing their diagrams side by side, whereas compositions of morphisms are depicted by putting one on top of the other, for instance the object $A \otimes B$, and the morphisms $f \otimes g$ and $f \circ h$, for $f\colon A \to B, g\colon C \to D$, and $h\colon B \to C$ are depicted as follows:

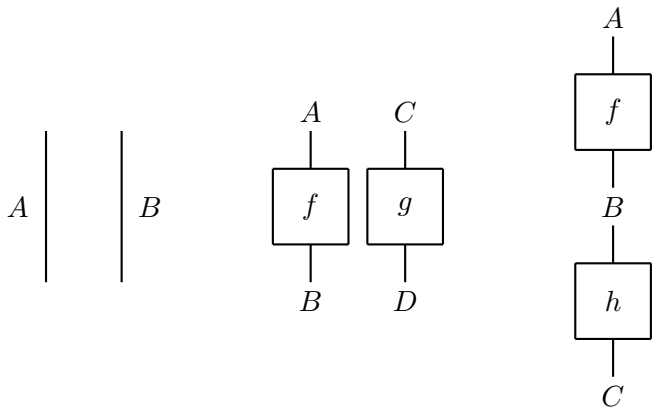

The ϵ maps are depicted by cups, η maps by caps, and yanking by their composition and straightening the strings. For instance, the diagrams for $\epsilon^l\colon A^l \otimes A \to I$, $\eta\colon I \to A \otimes A^l$ and $(\epsilon^l \otimes 1_A) \circ (1_A \otimes \eta^l) = 1_A$ are as follows:

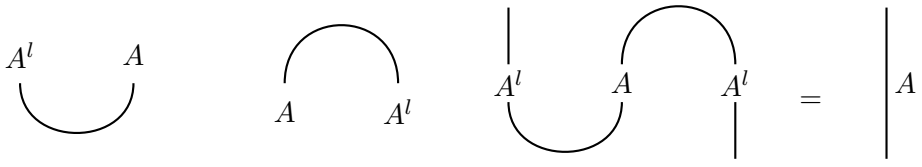

The composition of the ϵ and η maps with other morphisms is depicted as before, that is by juxtaposing them one above the other. For instance the compositions $(1_{B^l} \otimes f) \circ \epsilon^l$ and $\eta^l \circ (1_{A^l} \otimes f)$ are depicted as follows:

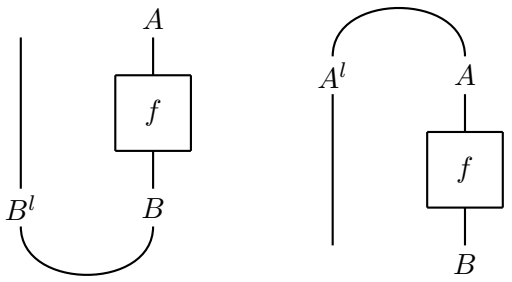

As for Frobenius algebras, the comonoid and monoid morphisms are depicted as follows:

The *Frobenius* condition is depicted as follows:

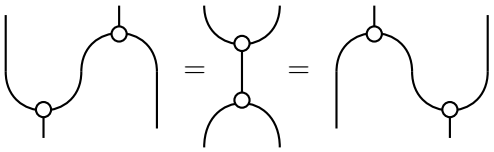

The defining equations of a commutative special Frobenius algebra guarantee that any picture depicting a Frobnius computation can be reduced to a normal form that only depends on the number of input and output strings of the nodes, independently of their topology. These normal forms can be simplified to so-called 'spiders':

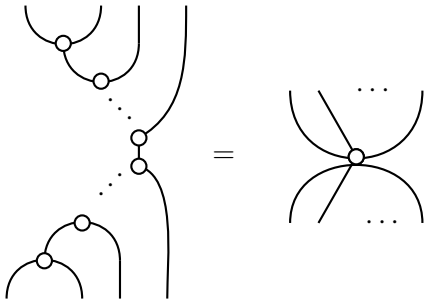

In the category $FVect$, apart from spaces V, W, which are objects of the category, we also have vectors $\overrightarrow{v}, \overrightarrow{w}$. These are depicted by their representing

morphisms and as triangles with a number of strings emanating from them. The number of strings of a triangle denote the tensor order of the vector, for instance, $\overrightarrow{v} \in V$, $\overrightarrow{v'} \in V \otimes W$, and $\overrightarrow{v''} \in V \otimes W \otimes Z$ are depicted as follows:

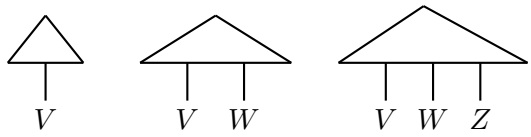

Application of an ϵ map to a vector is depicted similarly, for instance $\epsilon^l(\overrightarrow{v})$ is depicted by the composition $V^l \otimes V \xrightarrow{\epsilon^l} I \xrightarrow{\overrightarrow{v}} V$ and $\epsilon^r(\overrightarrow{v})$ by $V \otimes V^r \xrightarrow{\epsilon^r} I \xrightarrow{\overrightarrow{v}} V$. The diagrams for these applications are as follows:

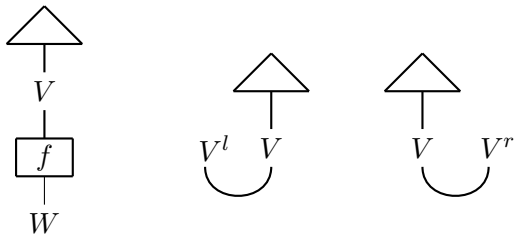

Applications of the Frobenius maps to vectors are depicted in a similar fashion; for instance $\mu(\overrightarrow{v} \otimes \overrightarrow{v})$ is the composition $I \otimes I \xrightarrow{\overrightarrow{v} \otimes \overrightarrow{v}} V \otimes V \xrightarrow{\mu} V$ and $\iota(\overrightarrow{v})$ is the composition $I \xrightarrow{\overrightarrow{v}} V \xrightarrow{\iota} I$, depicted as follows:

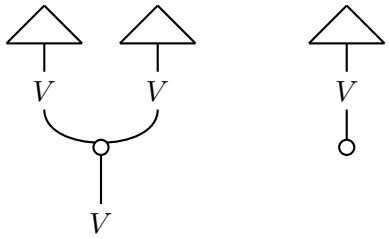

4 Vector Space Interpretations

The grammatical structure of a language is encoded in the category *Preg*: objects are grammatical types (assigned to words of the language) and morphisms are grammatical reductions (encoding the grammatical formation rules of the language). For instance, the grammatical structure of the sentence "Men love Mary" is encoded in the assignment of types n to the noun phrases "men" and "Mary" and $n^r \otimes s \otimes n^l$ to the verb "love", and in the reduction map $\epsilon^l_n \otimes 1_s \otimes \epsilon^r_n$. The application of this reduction map to the tensor product of the word types in the sentence results in the type s:

$$(\epsilon^l_n \otimes 1_s \otimes \epsilon^r_n)(n \otimes (n^r \otimes s \otimes n^l) \otimes n) \to s$$

To each reduction map corresponds a string diagram that depicts the structure of reduction:

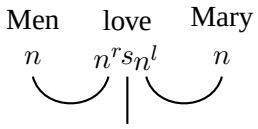

In Coecke et al. (2010) [8] the pregroup types and reductions are interpreted as vector spaces and linear maps, achieved via a homomorphic mapping from *Preg* to *FVect*. Categorically speaking, this map is a strongly monoidal functor:

$$F : \textit{Preg} \to \textit{FVect}$$

It assigns vector spaces to the basic types as follows:

$$F(1) = I \qquad F(n) = N \qquad F(s) = S$$

and to the compound types by monoidality as follows; for x, y objects of *Preg*:

$$F(x \otimes y) = F(x) \otimes F(y)$$

Monoidal functors preserve the compact structure; that is the following holds:

$$F(x^l) = F(x^r) = F(x)^*$$

For instance, the interpretation of a transitive verb is computed as follows:

$$F(n^r \otimes s \otimes n^l) = F(n^r) \otimes F(s) \otimes F(n^l) = \\ F(n)^* \otimes F(s) \otimes F(n)^* = N \otimes S \otimes N$$

This interpretation means that the meaning vector of a transitive verb is a vector in $N \otimes S \otimes N$.

The pregroup reductions, i.e. the partial order morphisms of *Preg*, are interpreted as linear maps: whenever $p \leq q$ in *Preg*, we have a linear map $f_{\leq} : F(p) \to F(q)$. The ϵ and η maps of *Preg* are interpreted as the ϵ and η maps of *FVect*. For instance, the pregroup reduction of a transitive verb sentence is computed as follows:

$$F(\epsilon_n^r \otimes 1_s \otimes \epsilon_n^r) = F(\epsilon_n^r) \otimes F(1_s) \otimes F(\epsilon_n^l) = \\ F(\epsilon_n)^* \otimes F(1_s) \otimes F(\epsilon_n)^* = \epsilon_N \otimes 1_S \otimes \epsilon_N$$

The distributional meaning of a sentence is obtained by applying the interpretation of the pregroup reduction of the sentence to the tensor product of the

distributional meanings of the words in the sentence. For instance, the distributional meaning of "Men love Mary" is as follows:

$$F(\epsilon_n^r \otimes 1_s \otimes \epsilon_n^l)(\overrightarrow{Men} \otimes \overrightarrow{love} \otimes \overrightarrow{Mary})$$

This meaning is depictable via the following string diagram:

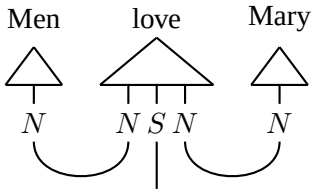

Symbolically, the above diagram corresponds to $(\epsilon_N \otimes 1_S \otimes \epsilon_N)(\overrightarrow{Men} \otimes \overrightarrow{love} \otimes \overrightarrow{Mary})$, which is the application of the interpretation of the pregroup reduction of the sentence to the tensor products of the vectors of the words within it. Depending on their types, the distributional meanings of the words are either atomic vectors or linear maps. For instance, the distributional meaning of 'men' is a vector in N, whereas the distributional meaning of 'love' is in the space $N \otimes S \otimes N$, hence a linear map, in this case from $N \otimes N$ to S.

The next section applies these techniques to the distributional interpretation of pronouns. The interpretations are defined using: ϵ maps, for application of the semantics of one word to another; η maps, to pass information around by bridging intermediate words; and Frobenius operations, for copying and combining the noun vectors and discarding the sentence vectors.

5 Modelling Relative Pronouns

In this paper we focus on the subject and object relative pronouns, *who(m)*, *which* and *that*. Examples of noun phrases with subject relative pronouns are "men who love Mary", "dog which ate cats". Examples of noun phrases with object relative pronouns are "men whom Mary loves", "book that John read". In the final example, "book" is the head noun, modified by the relative clause "that John read". The intuition behind the use of Frobenius algebras to model such cases is the following. In "book that John read", the relative clause acts on the noun (modifies it) via the relative pronoun, which passes information from the clause to the noun. The relative clause is then discarded, and the modified noun is returned. Frobenius algebras provide the machinery for all of these operations. The pregroup types of the relative pronouns are as follows:

$$n^r n s^l n \quad \text{(subject)} \qquad\qquad n^r n n^{ll} s^l \text{(object)}$$

As argued in Lambek (2008) [19], the occurrence of double left adjoints on the type of the object relative pronoun is a sign of the movement of the object from the end of the clause to its beginning. These types result in the following reductions:

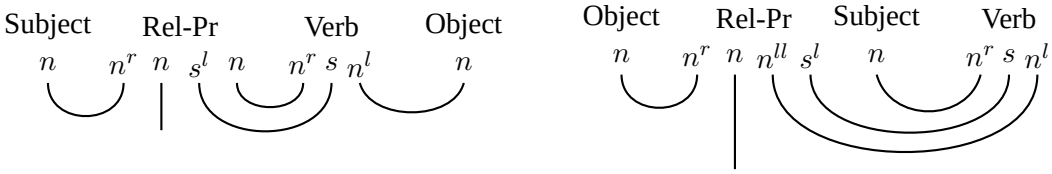

The meaning spaces of these pronouns are computed using the mechanism described above:

$$\begin{aligned} F(n^r n s^l n) &= F(n^r) \otimes F(n) \otimes F(s^l) \otimes F(n) \\ &= N \otimes N \otimes S \otimes N \\ F(n^r n n^{ll} s^l) &= F(n^r) \otimes F(n) \otimes F(n^{ll}) \otimes F(s^l) \\ &= N \otimes N \otimes N \otimes S \end{aligned}$$

The semantic roles that these pronouns play are reflected in their categorical vector space meanings, depicted as follows:

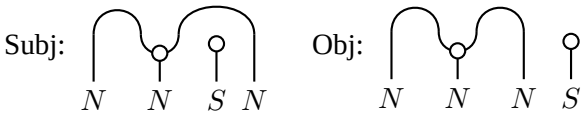

with the following corresponding morphisms:

$$\text{Subj: } (1_N \otimes \mu_N \otimes \zeta_S \otimes 1_N) \circ (\eta_N \otimes \eta_N)$$

$$\text{Obj: } (1_N \otimes \mu_N \otimes 1_N \otimes \zeta_S) \circ (\eta_N \otimes \eta_N)$$

The details of these compositions are:

$$\text{Subj: } \quad I \cong I \otimes I \xrightarrow{\eta_N \otimes \eta_N} N \otimes N \otimes N \otimes N \cong N \otimes N \otimes N \otimes I \otimes N \xrightarrow{1_N \otimes \mu_N \otimes \zeta_S \otimes 1_N} N \otimes N \otimes S \otimes N$$

$$\text{Obj: } \quad I \cong I \otimes I \xrightarrow{\eta_N \otimes \eta_N} N \otimes N \otimes N \otimes N \cong N \otimes N \otimes N \otimes N \otimes I \xrightarrow{1_N \otimes \mu_N \otimes 1_N \otimes \zeta_S} N \otimes N \otimes N \otimes S$$

The diagram of the meaning vector of the subject relative clause interacting with the head noun is as follows:

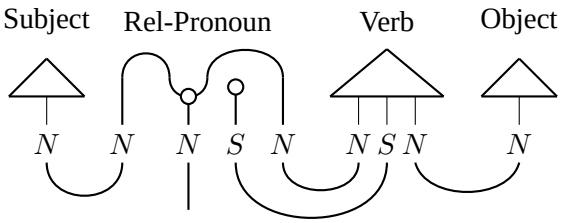

The diagram for the object relative clause is:

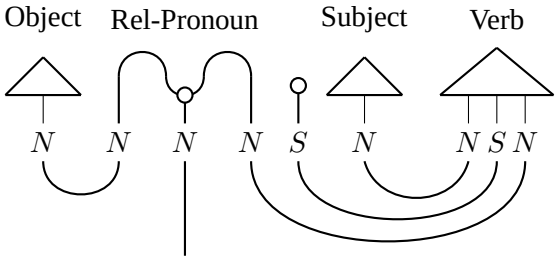

These diagrams depict the flow of information in a relative clause and the semantic role of its relative pronoun, which 1) passes information from the clause to the head noun via the η maps; 2) acts on the noun via the μ map; 3) discards the clause via the ζ map; and 4) returns the modified noun via 1_N. The ϵ maps pass the information of the subject and object nouns to the verb and to the relative pronoun to be acted on. Note that there are two different flows of information in these clauses: the ones that come from the grammatical structure and are depicted by ϵ maps (at the bottom of the diagrams), and the ones that come from the semantic role of the pronoun and are depicted by η maps (at the top of the diagrams).

The normal forms of these diagrams are:

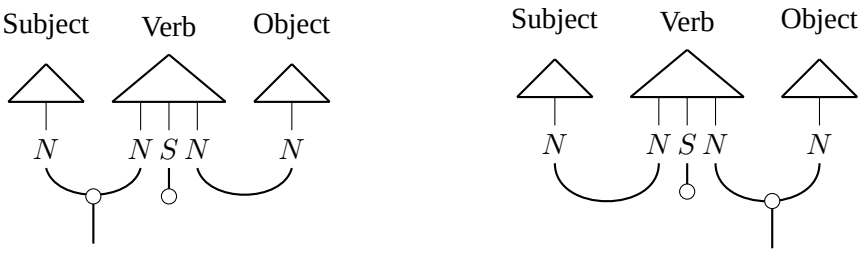

Symbolically, they correspond to the following morphisms:

$$(\mu_N \otimes \iota_S \otimes \epsilon_N)\left(\overrightarrow{\text{Subject}} \otimes \overrightarrow{\text{Verb}} \otimes \overrightarrow{\text{Object}}\right)$$
$$(\epsilon_N \otimes \iota_S \otimes \mu_N)\left(\overrightarrow{\text{Subject}} \otimes \overrightarrow{\text{Verb}} \otimes \overrightarrow{\text{Object}}\right)$$

The simplified normal forms will become useful in practice when calculating vectors for such cases.

6 Truth Theoretic Instantiations

In this section we demonstrate the effect of the Frobenius operations in a truth-theoretic setting, similar to Coecke et al. (2010) [8] but also allowing for degrees

of truth. This instantiation is designed only as a theoretical example, which merely recasts Montague semantics, rather than a suggestion for implementation. A suggestion for concrete implementation is presented in Section 8.

Take N to be the vector space spanned by a set of individuals $\{\overrightarrow{n}_i\}_i$ that are mutually orthogonal. For example, $\overrightarrow{n}_1$ represents the individual Mary, $\overrightarrow{n}_{25}$ represents Roger the dog, $\overrightarrow{n}_{10}$ represents John, and so on. A sum of basis vectors in this space represents a common noun; e.g. $\overrightarrow{man} = \sum_i \overrightarrow{n}_i$, where i ranges over the basis vectors denoting men. We take S to be the one dimensional space spanned by the single vector $\overrightarrow{1}$. The unit vector spanning S represents truth value 1, the zero vector represents truth value 0, and the intermediate vectors represent degrees of truth.

A transitive verb w, which is a vector in the space $N \otimes S \otimes N$, is represented as follows:

$$\overline{w} := \sum_{ij} \overrightarrow{n}_i \otimes (\alpha_{ij} \overrightarrow{1}) \otimes \overrightarrow{n}_j$$

if $\overrightarrow{n}_i$ w's $\overrightarrow{n}_j$ with degree α_{ij}, for all i, j.

Further, since S is one-dimensional with its only basis vector being $\overrightarrow{1}$, the transitive verb can be represented by the following element of $N \otimes N$:

$$\sum_{ij} \alpha_{ij} \overrightarrow{n}_i \otimes \overrightarrow{n}_j \quad \text{if} \quad \overrightarrow{n}_i \; w\text{'s } \overrightarrow{n}_j \text{ with degree } \alpha_{ij}$$

Restricting to either $\alpha_{ij} = 1$ or $\alpha_{ij} = 0$ provides a 0/1 meaning, i.e. either $\overrightarrow{n}_i$ w's $\overrightarrow{n}_j$ or not. Letting α_{ij} range over the interval $[0, 1]$ enables us to represent degrees as well as limiting cases of truth and falsity. For example, the verb "love", denoted by $\overline{love}$, is represented by:

$$\sum_{ij} \alpha_{ij} \overrightarrow{n}_i \otimes \overrightarrow{n}_j \quad \text{if} \quad \overrightarrow{n}_i \text{ loves } \overrightarrow{n}_j \text{with degree } \alpha_{ij}$$

If we take α_{ij} to be 1 or 0, from the above we obtain the following:

$$\sum_{(i,j) \in R_{love}} \overrightarrow{n}_i \otimes \overrightarrow{n}_j$$

where R_{love} is the set of all pairs (i, j) such that $\overrightarrow{n}_i$ loves $\overrightarrow{n}_j$. Note that, with this definition, the sentence space has already been discarded, and so for this instantiation the ι map, which is the part of the relative pronoun interpretation designed to discard the relative clause after it has acted on the head noun, is not required.

For common nouns $\overrightarrow{\text{Subject}} = \sum_{k \in K} \overrightarrow{n}_k$ and $\overrightarrow{\text{Object}} = \sum_{l \in L} \overrightarrow{n}_l$, where k and l range over the sets of basis vectors representing the respective common nouns, the truth-theoretic meaning of a noun phrase modified by a subject relative clause is computed as follows:

$$\begin{aligned}
&\overrightarrow{\text{Subject who Verb Object}} := (\mu_N \otimes \epsilon_N)\left(\overrightarrow{\text{Subject}} \otimes \overrightarrow{\text{Verb}} \otimes \overrightarrow{\text{Object}}\right) \\
&= (\mu_N \otimes \epsilon_N)\left(\sum_{k \in K} \overrightarrow{n}_k \otimes (\sum_{ij} \alpha_{ij} \overrightarrow{n}_i \otimes \overrightarrow{n}_j) \otimes \sum_{l \in L} \overrightarrow{n}_l\right) \\
&= \sum_{ij,k \in K, l \in L} \alpha_{ij} \mu_N(\overrightarrow{n}_k \otimes \overrightarrow{n}_i) \otimes \epsilon_N(\overrightarrow{n}_j \otimes \overrightarrow{n}_l) \\
&= \sum_{ij,k \in K, l \in L} \alpha_{ij} \delta_{ki} \overrightarrow{n}_i \delta_{jl} \\
&= \sum_{k \in K, l \in L} \alpha_{kl} \overrightarrow{n}_k
\end{aligned}$$

The result is highly intuitive, namely the sum of the subject individuals weighted by the degree with which they have acted on the object individuals via the verb. A similar computation, with the difference that the μ and ϵ maps are swapped, provides the truth-theoretic semantics of the object relative clause:

$$\sum_{k \in K, l \in L} \alpha_{kl} \overrightarrow{n}_l$$

The calculation and final outcome is best understood with an example.

Now only consider truth values 0 and 1. Consider the noun phrase with object relative clause "men whom Mary loves" and take N to be the vector space spanned by the set of all people; then the males form a subspace of this space, where the basis vectors of this subspace, i.e. men, are denoted by $\overrightarrow{m}_l$, where l ranges over the set of men which we denote by M. We set "Mary" to be the individual $\overrightarrow{f}_1$, "men" to be the common noun $\sum_{l \in M} \overrightarrow{m}_l$, and "love" to be as follows:

$$\sum_{(i,j) \in R_{love}} \overrightarrow{f}_i \otimes \overrightarrow{m}_j$$

The vector corresponding to the meaning of "men whom Mary loves" is computed as follows and as expected, the result is the sum of the men basis vectors which are also loved by Mary.

$$\overrightarrow{\text{men whom Mary loves}} := (\epsilon_N \otimes \mu_N)\left(\overrightarrow{f}_1 \otimes (\sum_{(i,j)\in R_{love}} \overrightarrow{f}_i \otimes \overrightarrow{m}_j) \otimes \sum_{l\in M} \overrightarrow{m}_l\right)$$

$$= \sum_{l\in M,(i,j)\in R_{love}} \epsilon_N(\overrightarrow{f}_1 \otimes \overrightarrow{f}_i) \otimes \mu(\overrightarrow{m}_j \otimes \overrightarrow{m}_l)$$

$$= \sum_{l\in M,(i,j)\in R_{love}} \delta_{1i}\delta_{jl}\overrightarrow{m}_j$$

$$= \sum_{(1,j)\in R_{love}|j\in M} \overrightarrow{m}_j$$

A second example involves degrees of truth. Suppose we have two females Mary $\overrightarrow{f}_1$ and Jane $\overrightarrow{f}_2$ and four men $\overrightarrow{m}_1, \overrightarrow{m}_2, \overrightarrow{m}_3, \overrightarrow{m}_4$. Mary loves $\overrightarrow{m}_1$ with degree 1/4 and $\overrightarrow{m}_2$ with degree 1/2; Jane loves $\overrightarrow{m}_3$ with degree 1/5; and $\overrightarrow{m}_4$ is not loved. In this situation, we have:

$$R_{love} = \{(1,1),(1,2),(2,3)\}$$

and the verb love is represented by:

$$1/4(\overrightarrow{f}_1 \otimes \overrightarrow{m}_1) + 1/2(\overrightarrow{f}_1 \otimes \overrightarrow{m}_2) + 1/5(\overrightarrow{f}_2 \otimes \overrightarrow{m}_3)$$

The meaning of "men whom Mary loves" is computed by substituting an $\alpha_{1,j}$ in the last line of the above computation, resulting in the men whom Mary loves together with the degrees that she loves them:

$$\sum_{(1,j)\in R_{love}|j\in M} \alpha_{1j}\overrightarrow{m}_j = 1/4\overrightarrow{m}_1 + 1/2\overrightarrow{m}_2$$

The meaning of the clause "men whom women love" is computed as follows, where W is the set of women:

$$\sum_{k\in W,l\in M,(i,j)\in R_{love}} \alpha_{ij}\epsilon_N(\overrightarrow{f}_k \otimes \overrightarrow{f}_i) \otimes \mu(\overrightarrow{m}_j \otimes \overrightarrow{m}_l)$$

$$= \sum_{k\in W,l\in M,(i,j)\in R_{love}} \alpha_{ij}\delta_{ki}\delta_{jl}\overrightarrow{m}_j$$

$$= \sum_{(i,j)\in R_{love}|i\in W,j\in M} \alpha_{ij}\overrightarrow{m}_j$$

$$= 1/4\overrightarrow{m}_1 + 1/2\overrightarrow{m}_2 + 1/5\overrightarrow{m}_3$$

The result is the men loved by Mary or Jane together with the degrees to which they are loved.

As a more complicated example, consider the clause ”Movies that Mary liked which became famous”, the first pronoun ‘that’ has an objective role and the second pronoun ‘which’ has a subjective role. The diagram of this clause is as follows:

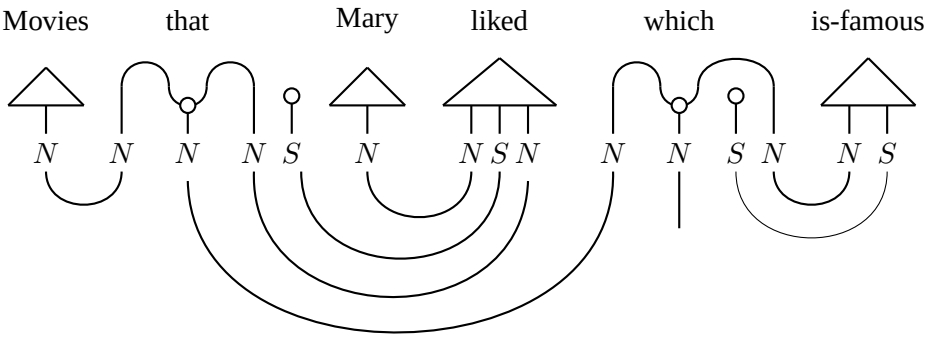

This first normalises to the following:

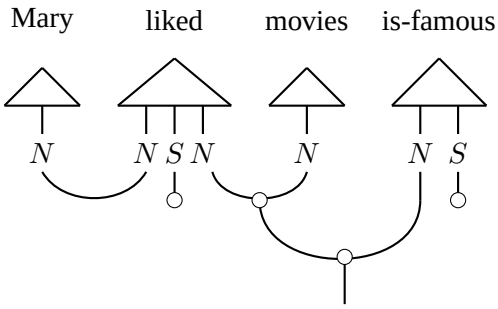

and then by the spider normal form, the two μ maps fuse and we obtain the following normal form:

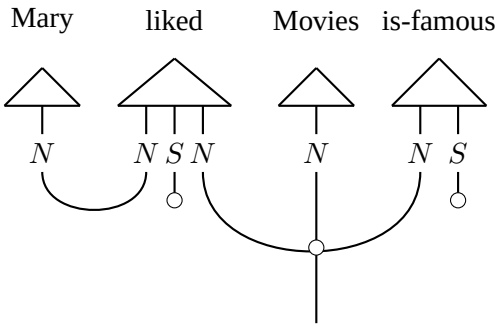

The vector space meaning of this clause is computed according to the first normalisation above (that is, before the application of the spider normal form) to demonstrate how the composition of two μ maps works (we also drop the ι map as our concrete truth-theoretic construction does not need it). The result is a weighted sum of the famous movies that Mary loves, where each weight is the multiplications of the degree to which Mary loves a movie with the degree to which the movie is famous.

We may also have cases where the two relative pronouns do not compose, for instance in ”men who love books that Mary wrote”. The diagram for this sentence is as follows:

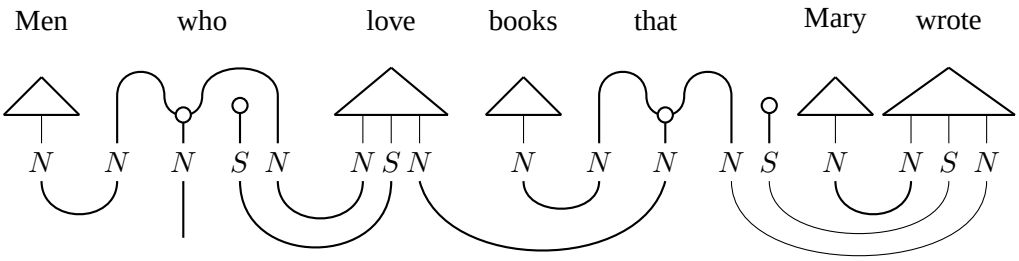

Here the μ maps do not compose: the first μ is applied to the subject of ”loves”, whereas the second one is applied to the object of ”wrote” and these two dimensions do not interact. The following simplified form of the above diagram makes this point clearer:

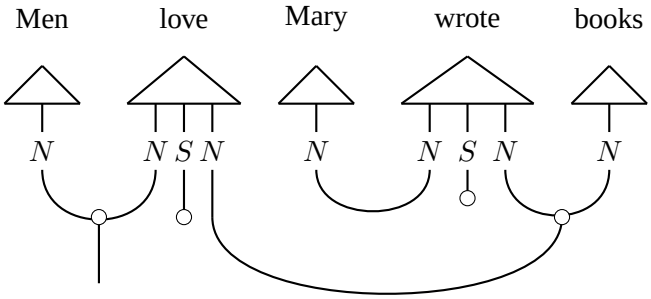

Here, the existing μ and ι's do not interact either: the first pair filters out men who love books from the set of all men and the second pair filters out books that Mary wrote from the set of all books. In the Boolean case, the result will be the sum of men who love books written by Mary. In the case of degrees of love, the result will be a weighted sum of men, each with the degree to which they love these books.

A a final example, consider the sentence ‘Mary loves men whom Mary loves’, which is a tautology, i.e. it is always true. The pictorial meaning of this sentence is as follows:

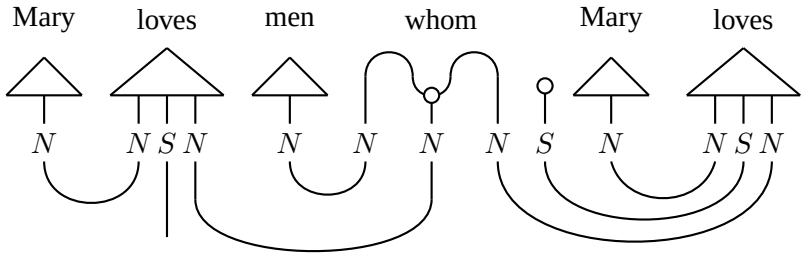

which normalises to

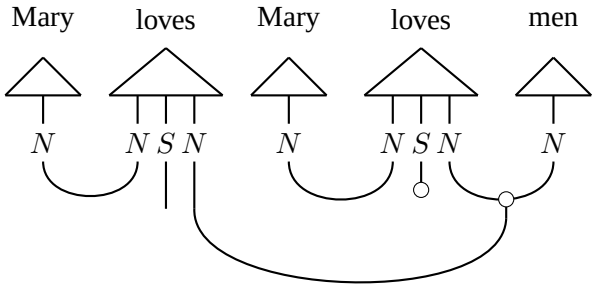

The vector space meaning of this sentence becomes as follows:

$$\sum_{ij} \langle \overrightarrow{f}_1 \mid \overrightarrow{f}_i \rangle \, \langle \overrightarrow{m}_j \mid \sum_{ij} \delta_{1i} \, \delta_{lj} \, \overrightarrow{m}_j \rangle$$

where $\sum_{ij} \delta_{1i} \, \delta_{lj} \, \overrightarrow{m}_j$ is ‘men whom Mary loves’, as computed above. So the inner product $\langle \overrightarrow{m}_j \mid \sum_{ij} \delta_{1i} \, \delta_{lj} \, \overrightarrow{m}_j \rangle$ only returns the men for which ’love Mary’ is true, resulting aways in the output 1.

7 Embedding Predicate Semantics

In this section, we provide a set-theoretic interpretation for relative clauses according to the constructions discussed in [26]. We start by fixing a universe of elements $\mathcal{U}$. A proper nouns is an individual (i.e. an element) in this set; a common noun is the set of the individuals that have the property expressed by the common noun; hence common nouns are unary predicates over $\mathcal{U}$. An intransitive verb is the set of all individuals that are acted upon by the relationship expressed by the verb, hence these are unary predicates over $\mathcal{U}$. A transitive verb is the set of pairs of individuals that are related by the relationship expressed by

the verb; hence these are binary predicates over $\mathcal{U} \times \mathcal{U}$. Here is an example for each case:

$$\begin{aligned}
[\![Mary]\!] &= \{u_1\} \\
[\![Author]\!] &= \{x \in \mathcal{U} \mid \text{"}x \text{ is an author"}\} \\
[\![Sleep]\!] &= \{x \in \mathcal{U} \mid \text{"}x \text{ sleeps"}\} \\
[\![Entertain]\!] &= \{(x, y) \in \mathcal{U} \times \mathcal{U} \mid \text{"}x \text{ entertains } y\text{"}\}
\end{aligned}$$

The subjective and objective relative clauses are interpreted as follows:

$$\begin{aligned}
&\{x \in \mathcal{U} \mid x \in [\![Subj]\!], (x, y) \in [\![Verb]\!], \text{ for all } y \in [\![Obj]\!]\} \\
&\{y \in \mathcal{U} \mid y \in [\![Obj]\!], (x, y) \in [\![Verb]\!], \text{ for all } x \in [\![Subj]\!]\}
\end{aligned}$$

These clauses pick out individuals which belong to the interpretation of subject/object and which are in the relationship expressed by the verb with all elements of the interpretation of object/subject. Examples are 'men who love Mary' and 'men whom cats loves', interpreted as follows, for $x \in \mathcal{U}$:

$$\begin{aligned}
&\{x \in \mathcal{U} \mid x \in [\![men]\!], (x, y) \in [\![Love]\!], \text{ for all } y \in [\![Mary]\!]\} \\
&\{y \in \mathcal{U} \mid y \in [\![men]\!], (x, y) \in [\![Love]\!], \text{ for all } x \in [\![cat]\!]\}
\end{aligned}$$

The first example picks out individuals that are 'men' and who are in relationship of 'love' with the individual 'Mary'. The second one picks out individuals that are 'men' and who are 'loved' by all the individual that are ' cats'.

We obtain the vector space forms of the above predicate interpretations by developing a map from the category of sets and relations to the category of finite dimensional vector spaces with a fixed orthogonal basis and linear maps over them, that is a map with types as follows:

$$e \colon Rel \hookrightarrow FVect$$

We set this map to turn a set T into a vector space V_T spanned by the elements of T and an element $t \in T$ into a basis vector $\overrightarrow{t}$ of V_T. A subset W of T is turned into the sum vector of its elements, i.e. $\sum_i \overrightarrow{w}_i$, where w_i's enumerate over elements of W. A relation $R \subseteq T \times T'$ is turned into a linear map from T to T', or equivalently as an element of the space $T \otimes T'$; we represent this element by the sum of its basis vectors $\sum_{ij} \overrightarrow{t}_i \otimes \overrightarrow{t'}_j$, here t_i's enumerate elements of T and t'_j's enumerate elements of T'.

We then use the compact structure of Rel and $FVect$ and turn the application of a relation $R \subseteq T \times T'$ to its arguments into the inner product of the

vector space forms of the relation and the arguments. In our sum representation, $R(a)$ when $a \in T$ and $R^{-1}(b)$ when $b \in T'$ are turned into the following:

$$\sum_{ij} \langle a \mid \overrightarrow{t}_i \rangle \overrightarrow{t'}_j \qquad\qquad \sum_{ij} \overrightarrow{t}_i \langle \overrightarrow{t'}_j \mid b \rangle$$

To check wether an element t is in the set T, we use the Frobenius operation μ on T, that is $\mu \colon T \otimes T \to T$, as follows:

$$t \in T \quad \text{whenever} \quad \mu(\overrightarrow{t}, \sum_i \overrightarrow{t}_i) = \overrightarrow{t}$$

The intersection of two sets T, T' is computed by generalising the above via the Frobenius μ operation on the whole universe, that is $\mu \colon \mathcal{U} \otimes \mathcal{U} \to \mathcal{U}$, or on a set containing both of these sets, e.g. $T \cup T'$, that is $\mu \colon (T \cup T') \otimes (T \cup T') \to (T \cup T')$. In either case, for $\sum_i \overrightarrow{t}_i$ and $\sum_j \overrightarrow{t'}_j$ representations of T and T' we obtain the following for the intersection:

$$T \cap T' \;:=\; \mu(\sum_i \overrightarrow{t}_i, \sum_j \overrightarrow{t'}_j)$$

We are now ready to show that the vector space forms of the predicate interpretation of relative clauses, developed along side the constructions described above, provide us with the same semantics as the truth-theoretic vector space semantics developed in Section 6. That is, we have:

Proposition 1. *The map e, described above, provides a 0/1 truth-theoretic interpretation for relative clauses.*

Proof. The first step is to instantiate the proper and common nouns, the intransitive and transitive verbs. For this, we set the universe $\mathcal{U}$ to be the individuals representing the nouns of the language, then the map e instantiates as follows:

- The universe $\mathcal{U}$ becomes the vector space $V_{\mathcal{U}}$, spanned by its elements. So an element $u_i \in \mathcal{U}$ becomes a basis vector $\overrightarrow{u}_i$ of $V_{\mathcal{U}}$.
- A proper noun $a \in \mathcal{U}$ becomes a basis vector $\overrightarrow{u}_a$ of $V_{\mathcal{U}}$.
- A common noun $c \subset \mathcal{U}$ becomes a subspace of $V_{\mathcal{U}}$, represented by the sum of the representations of its elements $\overrightarrow{u}_i$, that is

$$c \hookrightarrow \overrightarrow{c} \quad :: \quad \sum_i \overrightarrow{u}_i \quad \text{for all } u_i \in c$$

- An intransitive verb $v \subset \mathcal{U}$ becomes a subspace of $V_{\mathcal{U}}$, represented by the sum of all its elements u_i, that is

$$v \hookrightarrow \overrightarrow{v} \quad :: \quad \sum_i \overrightarrow{u}_i \quad \text{for all } v_i \in v$$

– A transitive verb $w \subset \mathcal{U} \times \mathcal{U}$ becomes a linear map $\overline{w}$ in $V_{\mathcal{U}} \otimes V_{\mathcal{U}}$, represented by the sum of its basis vectors $\overrightarrow{u}_i, \overrightarrow{u}_j$, which in this case are pairs of its elements (u_i, u_j), that is

$$w \hookrightarrow \overline{w} \quad :: \quad \sum_{ij} \overrightarrow{u}_i \otimes \overrightarrow{u}_j \quad \text{for all } (u_i, u_j) \in w$$

The next step is to use the above definitions and develop the vector space forms of the predicate interpretations of the relative clauses. To do so, we turn these interpretations into intersections of subsets and check the membership relation via the μ map. As for the application of the predicates to the interpretation of their arguments, we use their relational image on subsets. That is, for $R \subseteq \mathcal{U} \times \mathcal{U}$ and $T \subseteq \mathcal{U}$, we work with $R[T]$, defined by $R[T] = \{t' \in \mathcal{U} \mid (t, t') \in R \text{ for } t \in T\}$. This form of application has an implicit universal quantification: it consists of all the elements of $\mathcal{U}$ that are related to some element of T. Hence, the intersection forms take care of the quantification present in the predicate semantics of the relative clauses without having to explicitly use it.

Consider the subjective clause, its predicate interpretation is equivalent to the following intersection of subsets:

$$[\![Subj]\!] \cap [\![Verb]\!]^{-1}[[\![Obj]\!]]$$

The vector space form of this intersection is the application of the μ map to the vector space forms of each of the subsets involved in it. For the latter, we use their sum representation. That is, for $\sum_i \overrightarrow{u}_i$ and $\sum_j \overrightarrow{u}_j$ the representation of subspaces $V_{[\![Subj]\!]}$ and $V_{[\![Verb]\!]^{-1}[[\![Obj]\!]]}$, the vector space form of the above set is obtained as follows:

$$V_{[\![Subj]\!] \cap [\![Verb]\!]^{-1}[[\![Obj]\!]]} = \mu\left(V_{[\![Subj]\!]}, V_{[\![Verb]\!]^{-1}[[\![Obj]\!]]}\right) = \mu\left(\sum_i \overrightarrow{u}_i, \sum_j \overrightarrow{u}_j\right)$$

Recall that relational application was modelled by taking inner products, so we obtain $\sum_j \overrightarrow{u}_j$ by taking the inner product of vector space form of the verb to the vector form of the object, that is we have:

$$V_{[\![Verb]\!]^{-1}[[\![Obj]\!]]} = \sum_j \overrightarrow{u}_j = \sum_{ij} \overrightarrow{u}_i \langle \overrightarrow{u}_j \mid \sum_{b \in V_{[\![Obj]\!]}} \overrightarrow{u}_b \rangle$$

Substituting this term for $\sum_j \overrightarrow{u}_j$ in the μ term above provides us with the truth-theoretic meaning of the same clause in vector spaces. For the other clauses, we follow a similar procedure on the intersection form of their predicate semantics. The case for the objective relative clause, whose intersection form is $[\![Obj]\!] \cap [\![Verb]\!][[\![Subj]\!]]$ is almost immediate.

8 Concrete Instantiations

In the model of Grefenstette and Sadrzadeh (2011a) [10], the meaning of a verb is taken to be "the degree to which the verb relates properties of its subjects to properties of its objects". Clark (2013) [4] provides some examples showing how this is an intuitive definition for a transitive verb in the categorical framework. This degree is computed by forming the sum of the tensor products of the subjects and objects of the verb across a corpus, where w ranges over instances of the verb:

$$\overline{\text{verb}} = \sum_w (\overrightarrow{\text{sbj}} \otimes \overrightarrow{\text{obj}})_w$$

We denote the vector space of nouns by N; the above is a matrix in $N \otimes N$, depicted by a two-legged triangle as follows:

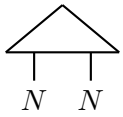

The verbs of this model do not have a sentence dimension; hence no information needs to be discarded when they are used in our setting, and so no ι map appears in the diagram of the relative clause. Inserting the above diagram in the diagrams of the normal forms results in the following for the subject relative clause (the object case is similar):

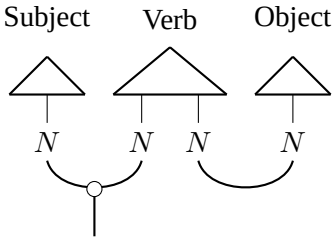

The abstract vectors corresponding to such diagrams are similar to the truth-theoretic case, with the difference that the vectors are populated from corpora and the scalar weights for noun vectors are not necessarily 1 or 0. For subject and object noun context vectors computed from a corpus as follows:

$$\overrightarrow{\text{Subject}} = \sum_k C_k \overrightarrow{n}_k \qquad \overrightarrow{\text{Object}} = \sum_l C_l \overrightarrow{n}_l$$

and the verb a linear map:

$$\overline{\text{Verb}} = \sum_{ij} C_{ij} \overrightarrow{n}_i \otimes \overrightarrow{n}_j$$

computed as above, the concrete meaning of a noun phrase modified by a subject relative clause is as follows:

$$\begin{aligned}
&\sum_{kijl} C_k C_{ij} C_l \mu_N(\overrightarrow{n}_k \otimes \overrightarrow{n}_i) \epsilon_N(\overrightarrow{n}_j \otimes \overrightarrow{n}_l) \\
&= \sum_{kijl} C_k C_{ij} C_l \delta_{ki} \overrightarrow{n}_k \delta_{jl} \\
&= \sum_{kl} C_k C_{kl} C_l \overrightarrow{n}_k
\end{aligned}$$

Comparing this to the truth-theoretic case, we see that the previous α_{kl} are now obtained from a corpus and instantiated to $C_k C_{kl} C_l$. To see how the above expression represents the meaning of the noun phrase, decompose it into the following:

$$\sum_k C_k \overrightarrow{n}_k \odot \sum_{kl} C_{kl} C_l \overrightarrow{n}_l$$

Note that the second term of the above, which is the application of the verb to the object, modifies the subject via point-wise multiplication. A similar result arises for the object relative clause case.

As an example, suppose that N has two dimensions with basis vectors $\overrightarrow{n}_1$ and $\overrightarrow{n}_2$, and consider the noun phrase "dog that bites men". Define the vectors of "dog" and "men" as follows:

$$\overrightarrow{\text{dog}} = d_1 \overrightarrow{n}_1 + d_2 \overrightarrow{n}_2 \qquad \overrightarrow{\text{men}} = m_1 \overrightarrow{n}_1 + m_2 \overrightarrow{n}_2$$

and the matrix of "bites" by:

$$b_{11} \overrightarrow{n}_1 \otimes \overrightarrow{n}_2 + b_{12} \overrightarrow{n}_1 \otimes \overrightarrow{n}_2 + b_{21} \overrightarrow{n}_2 \otimes \overrightarrow{n}_1 + b_{22} \overrightarrow{n}_2 \otimes \overrightarrow{n}_2$$

Then the meaning of the noun phrase becomes:

$$\begin{aligned}
&\overrightarrow{\text{dog that bites men}} := \\
&d_1 b_{11} m_1 \overrightarrow{n}_1 + d_1 b_{12} m_2 \overrightarrow{n}_1 + d_2 b_{21} m_1 \overrightarrow{n}_2 + d_2 b_{22} m_2 \overrightarrow{n}_2 \; = \\
&(d_1 \overrightarrow{n}_1 + d_2 \overrightarrow{n}_2) \odot ((b_{11} m_1 + b_{12} m_2) \overrightarrow{n}_1 + (b_{21} m_1 + b_{22} m_2) \overrightarrow{n}_2)
\end{aligned}$$

Using matrix notation, we can decompose the second term further, from which the application of the verb to the object becomes apparent:

$$\begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \times \begin{pmatrix} m_1 \\ m_2 \end{pmatrix}$$

Hence for the whole clause we obtain:

$$\overrightarrow{\text{dog}} \odot (\overline{\text{bites}} \times \overrightarrow{\text{men}})$$

Again this result is highly intuitive: assuming that the basis vectors of the noun space represent properties of nouns, the meaning of "dog that bites men" is a vector representing the properties of dogs, which have been modified (via multiplication) by those properties of individuals which bite men. Put another way, those properties of dogs which overlap with properties of biting things get accentuated.

In the same lines, one can also depict the model of [11], where a matrix is built for the verb by taking the tensor product of the context vector of the verb with itself. Assuming that the context vector of verb w is $\overrightarrow{w}$, this matrix is given and depicted as follows:

$$\overline{w} = \overrightarrow{w} \otimes \overrightarrow{w} \qquad \text{depicted by}$$

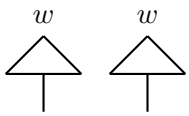

This instantiation blocks the flow of information from the subject to the object of the verb, and moreover either the application of the verb to the subject or to the object will be a scalar. The case of the subjective relative clause, where the application of the verb to its object becomes a scalar, is depicted bellows:

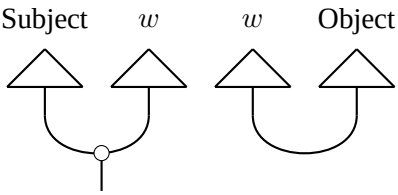

Most (if not all) of the distributional language tasks are based on the distance between the vectors with the cosine of the angle proven to be a good measure of word similarity. The cosine of the angle between two vectors is proportional to their inner product, hence it can be depicted using string diagrams, as follows:

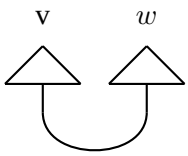

The problem with this model is that it will result in the same angle between, for instance, subjective clauses that have different objects, hence ignoring the object part all together (in the objective clauses the subject part will be ignored), as shown below:

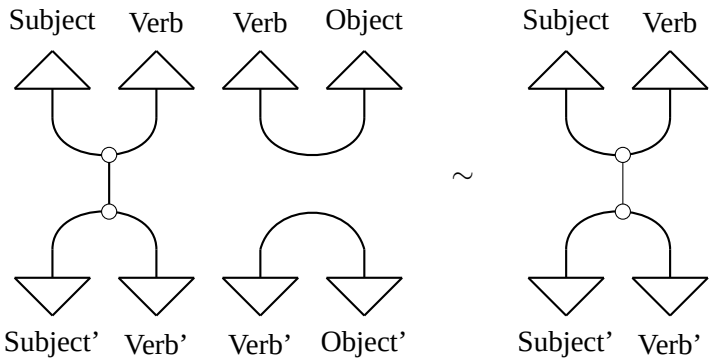

Here, the application of the verb to the object results in a scalar, hence the angle between ”Subj Verb Obj” and ”Subj’ Verb’ Obj’ ” will become equal to the angle between ”Subj Verb” and ”Subj’ Verb’ ”. A similar problem arises for other clauses.

Another possibility is the models of [13], which are obtained by starting from the model of [10], then either copying the subject or the object dimension of the verb, depicted as follows:

The problem here is that both of these models result in the same vector for relative clauses. To see why, consider the subjective relative clause, both models will results in the same vector, depicted as follows:

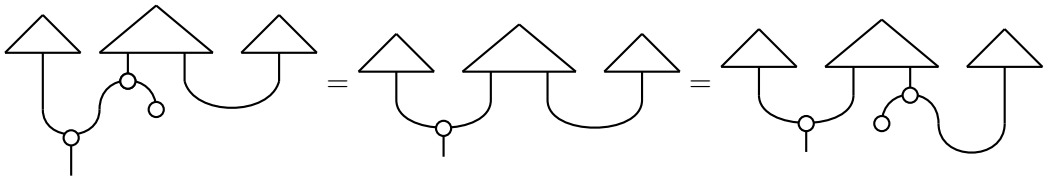

In the above models, one could substitute the concrete linear maps of verbs into the normal forms of relative clauses, as those models were formed based on a categorical setting and the type of the verb matched with the spaces required by our setting. This is not the case for the multiplicative model of [21], since here the meaning vector of a string of words $\mathrm{w}_1\mathrm{w}_2 \cdots \mathrm{w}_n$ is the component wise multiplication of their context vectors $\overrightarrow{w}_1 \odot \overrightarrow{w}_2 \odot \cdots \odot \overrightarrow{w}_n$. For instance, the meaning vector of a subjective relative clause ”Subj who Verb Obj” is the vector $\overrightarrow{\text{Subj}} \odot \overrightarrow{\text{who}} \odot \overrightarrow{\text{Verb}} \odot \overrightarrow{\text{Obj}}$.

In the compact categorical setting, component wise multiplication of two vectors can be modelled by the μ map. For two vectors $\overrightarrow{v} = \sum_i a_i \overrightarrow{x}_i$ and

$\overrightarrow{w} = \sum_j b_j \overrightarrow{x}_j$ we have:

$$\mu(\overrightarrow{v}, \overrightarrow{w}) \;=\; \mu(\overrightarrow{v} \otimes \overrightarrow{w}) \;=\; \sum_i a_i b_i \overrightarrow{x}_i = \overrightarrow{v} \odot \overrightarrow{w}$$

As a result we obtain $\overrightarrow{v} \odot \overrightarrow{w} = \overrightarrow{w} \odot \overrightarrow{v}$; diagrammatically we have:

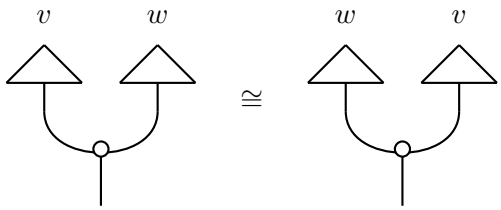

One problem with this model is the meaning of 'who' , which is either ignored or taken to be its context vector, both of which are problematic. A more serious problem is that this model does not respect the word order, hence cannot distinguish between grammatical and ungrammatical strings, and even in the grammatical cases, it assigns equal meanings to strings with the same words but in different orders. For instance, the clause 'men who eat rabbits' and 'rabbits who eat men' will have the same meaning, which will be the same as the vector for the ungrammatical strings 'rabbits men ate'.

9 A Toy Experiment

For demonstration purposes and in order to get an idea on how the data from a large scale corpus respond to the abstract computational methods developed here, we implement the concrete instantiation of Section 8 on the British National Corpus (BNC) and do a small-scale experiment, using parameters of previous experimentations with the categorical model[10,11,14,13].

The task we consider is a term/description classification task, similar to the term/definition classification task of Kartsaklis et al (2012) [13], where we replace *definitions* with *descriptions*. The motivation behind this task is the fact that relative clauses are often used to describe or provide extra information about words. For instance, the word 'football' may by described by the clause 'game that boys like' and the word 'doll' by the clause 'toy that girls prefer'. For our experiment, we chose a set of words and manually described each of them by an appropriate relative clause. This data set is presented in the following table:

	Term	Description
1	emperor	person who rule empire
2	queen	woman who reign country
3	mammal	animal which give birth
4	plug	plastic which stop water
5	carnivore	animal who eat meat
6	vegetarian	person who prefer vegetable
7	doll	toy that girl prefer
8	football	game that boy like
9	skirt	garment that woman wear

The goal of the task is to compute for what percentage of the words, their vectors are closest to the vectors of their descriptions. For this purpose, we compared the vector of each term to the vectors of all the descriptions. From our 9 terms, the cosines of 6 of them were the closest to their own description. The terms whose descriptions were not the closest to them were ‘plug, carnivore’, and ‘vegetarian’. In all of these cases the correct description was the second closest to the term. Below are some exemplary term/description cosines.

Term	Description	Cosine
football	game that boys like	0.62
	woman who reigns country	0.24
	toy that girls prefer	0.18
doll	toy that girls prefer	0.36
	woman who reigns country	0.35
	man who rules empire	0.30
mammal	animal which gives birth	0.36
	man who rules empire	0.19
	woman who reigns country	0.17
ball	man who rules empire	0.35
	woman who reigns country	0.28
	toy that girls prefer	0.22
plug	toy that girls prefer	0.24
	plastic which stops water	0.22
	woman who reigns country	0.17

To have a comparison baseline, we also experimented with a multiplicative and an additive model. In these models, one builds vectors for the relative clauses by multiplying/summing the vectors of the words in them in two ways:

(1) treating the relative pronoun as noise and not considering it in the computation, (2) treating the relative pronoun as any other word and considering its context vector in the computation. As prevue, the results were not to be significantly different (as the vectors of relative pronouns were dense and computing by them was similar to computing with the vector consisting of all 1's). For instance, the cosine between the vector of 'football' and its description was 0.39 when the pronoun was not considered and 0.37 when it was considered. Both of these are lower than in our model, where this cosine was 0.61. The first model got 6 out of the 9 cases correct; in the second model this number decreased to 5 out of 9. The individual results were mixed, in that they did bad on some of the good results of the Frobenius model, for example in the multiplicative model, the closest description to the term 'mammal' was 'animal that eats meat'. At the same time, they performed better on some of the bad terms of the Frobenious model, for example the description of 'plug' in the multiplicative model became the closest to it. However, the above numbers are not real representatives of the performance of the models; as the dataset is small and hand-made. Extending this task to a real one on a large automatically built dataset and doing more involved statistical analysis on the results requires a different venue; it constitutes work-in-progress.

10 Conclusion and Future Directions

In this paper, we have extended the compact categorical semantics of Coecke et al. (2010) [8] to analyse meanings of relative clauses in English from a vector space point of view. The resulting vector space semantics of the pronouns and clauses is based on the Frobenius algebraic operations on vector spaces: they reveal the internal structure, or what we call *anatomy*, of the relative clauses.

The methodology pursued in this paper and the Frobenius operations can be used to provide semantics for other relative pronouns and also other closed-class words such as prepositions and determiners. In each case, the grammatical type of the word and a detailed analysis of the role of these words in the meaning of the phrases in which they occur would be needed. In some cases, it may be necessary to introduce a linear map to represent the meaning of the word, for instance to distinguish the preposition *on* from *in*.

The contribution of this paper can be better demonstrated by using the string diagrammatic representations of the vector space meanings of these clauses. A noun phrase modified by a subject relative clause, which before this paper was depicted as follows:

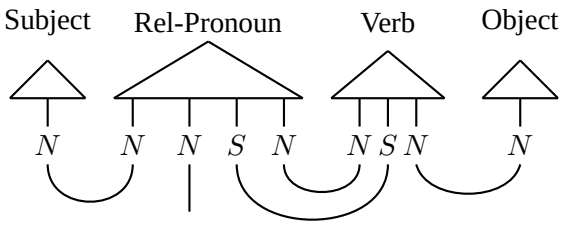

will now include the internal anatomy of its relative pronoun:

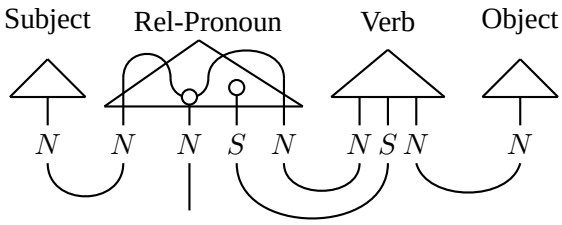

This internal structure shows how the information from the noun flows through the relative pronoun to the rest of the clause and how it interacts with other words. We have instantiated this vector space semantics using truth-theoretic and corpus-based examples.

One aspect of our example spaces which means that they work particularly well is that the sentence dimension in the verb is already discarded, which means that the ι maps are not required (as discussed above). Another feature is that the simple nature of the models means that the μ map does not lose any information, even though it takes the diagonal of a matrix and hence in general throws information away. The effect of the ι and μ maps in more complex representations of the verb remains to be studied in future work.

On the practical side, what we offer in this paper is a method for building appropriate vector representations for relative clauses. As a result, when presented with a relative clause, we are able to build a vector for it, only by relying on the vector representations of the words in the clause and the grammatical role of the relative pronoun. We do not need to retrieve information from a corpus to be able to build a vector or linear map for the relative pronoun, neither will we end up having to discard the pronoun and ignore the role that it plays in the meaning of the clause (which was perhaps the best option available before this paper). However, the Frobenius approach and our claim that the resulting vectors are 'appropriate' requires an empirical evaluation. Tasks such as the term definition task from Kartsaklis et al. (2010) [14] (which also uses Frobenius algebras but for a different purpose) are an obvious place to start. More generally, the sub-field of compositional distributional semantics is a growing and active

one [20,2,29,28], for which we argue that high-level mathematical investigations such as this paper, and also Clarke (2008) [5], can play a crucial role.

11 Acknowledgements

We would like to thank Dimitri Kartsaklis and Laura Rimell for helpful comments. Stephen Clark was supported by ERC Starting Grant DisCoTex (30692). Bob Coecke and Stephen Clark are supported by EPSRC Grant EP/I037512/1. Mehrnoosh Sadrzadeh is supported by an EPSRC CAF EP/J002607/1.

References

1. Baez, J., Dolan, J.: Higher-dimensional algebra and topological quantum field theory. Journal of Mathematical Physics 36, 6073–6105 (1995)
2. Baroni, M., Zamparelli, R.: Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In: Conference on Empirical Methods in Natural Language Processing (EMNLP-10). Cambridge, MA (2010)
3. Carboni, A., Walters, R.F.C.: Cartesian bicategories. I. J. Pure and Appied Algebra 49, 11–32 (1987)
4. Clark, S.: Type-driven syntax and semantics for composing meaning vectors. In: Heunen, C., Sadrzadeh, M., Grefenstette, E. (eds.) Quantum Physics and Linguistics: A Compositional, Diagrammatic Discourse, pp. 359–377. Oxford University Press (2013)
5. Clarke, D.: Context-theoretic Semantics for Natural Language: An Algebraic Framework. Ph.D. thesis, University of Sussex (2008)
6. Coecke, B., Paquette, E.: Introducing categories to the practicing physicist. In: Coecke, B. (ed.) New Structures for Physics, Lecture Notes in Physics, vol. 813, pp. 167–271. Springer (2008)
7. Coecke, B., Pavlovic, D., Vicary, J.: A new description of orthogonal bases. Mathematical Structures in Computer Science 1, 269–272 (2008)
8. Coecke, B., Sadrzadeh, M., Clark, S.: Mathematical foundations for a compositional distributional model of meaning. Linguistic Analysis 36, 345–384 (2010)
9. Frobenius, F.G.: Theorie der hyperkomplexen Größen. Preussische Akademie der Wissenschaften Berlin: Sitzungsberichte der Preußischen Akademie der Wissenschaften zu Berlin, Reichsdr (1903)
10. Grefenstette, E., Sadrzadeh, M.: Experimental support for a categorical compositional distributional model of meaning. In: Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 1394–1404 (2011)
11. Grefenstette, E., Sadrzadeh, M.: Experimenting with transitive verbs in a discocat. In: Proceedings of the Workshop on Geometrical Models of Natural Language Semantics (GEMS) (2011)
12. Joyal, A., Street, R.: The Geometry of Tensor Calculus, I. Advances in Mathematics 88, 55–112 (1991)
13. Kartsaklis, D., Sadrzadeh, M., Pulman, S.: A unified sentence space for categorical distributional-compositional semantics: Theory and experiments. In: Proceedings of the 24th International International Conference on Computational Linguistics (CoLing). pp. 193–213 (2012)

14. Kartsaklis, D., Sadrzadeh, M., Pulman, S., Coecke, B.: Compact Closed Categories and Frobenius Algebras for Natural Language Meaning. In: Chubb, J., Eskandarian, A., Harizanov, V. (eds.) Logic and Algebraic Structures in Quantum Computing and Information. Association for Symbolic Logic Lecture Notes in Logic, Cambridge University Press (2010)
15. Kelly, G.M., Laplaza, M.L.: Coherence for compact closed categories. Journal of Pure and Applied Algebra 19, 193–213 (1980)
16. Kock, J.: Frobenius algebras and 2D topological quantum field theories, London Mathematical Society student texts, vol. 59. Cambridge University Press (2003)
17. Lambek, J.: The Mathematics of Sentence Structure. American Mathematics Monthly 65, 154–170 (1958)
18. Lambek, J.: Type Grammar Revisited Logical Aspects of Computational Linguistics. In: Logical Aspects of Computational Linguistics, Lecture Notes in Computer Science, vol. 1582, pp. 1–27. Springer Berlin / Heidelberg (1999)
19. Lambek, J.: From Word to Sentence. Polimetrica / Milan (2008)
20. Mitchell, J., Lapata, M.: Vector-based models of semantic composition. In: Proceedings of ACL-08. pp. 236–244. Columbus, OH (2008)
21. Mitchell, J., Lapata, M.: Vector-based models of semantic composition. In: Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL). pp. 236–244 (2008)
22. Montague, R.: English as a formal language. In: Thomason, R.H. (ed.) Formal philosophy: Selected Papers of Richard Montague, pp. 189–223. Yale University Press (1974)
23. Moortgat, M.: Categorical type logics. In: van Benthem, J., ter Meulen, A. (eds.) Handbook of Logic and Language, pp. 93–177. Elsevier (1997)
24. Morrill, G.: Categorial Grammar. Oxford University Press (2012)
25. Preller, A., Lambek, J.: Free compact 2-categories. Mathematical Structures in Computer Science 17, 309–340 (2007)
26. Sag, I.A.: English relative clause constructions. Journal of Linguistics 33, 431–484 (1997)
27. Schütze, H.: Automatic word sense discrimination. Computational Linguistics 24, 97–123 (1998)
28. Socher, R., Pennington, J., Huang, E., Ng, A.Y., Manning, C.D.: Semi-supervised recursive autoencoders for predicting sentiment distributions. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. Edinburgh, UK (2011)
29. Zanzotto, F.M., Korkontzelos, I., Fallucchi, F., Manandhar, S.: Estimating linear models for compositional distributional semantics. In: Proceedings of COLING. Beijing, China (2010)